arXiv:1608.01793v1 [cs.CV] 5 Aug 2016

Sparse Subspace Clustering via Diffusion Process

Qilin Li, Ling Li, Wanquan Liu
Curtin University, Perth, Australia
kylinlovesummer@gmail.com

Abstract

Subspace clustering refers to the problem of clustering high-dimensional data that lie in a union of low-dimensional subspaces. State-of-the-art subspace clustering methods are based on the idea of expressing each data point as a linear combination of other data points while regularizing the matrix of coefficients with ℓ_1, ℓ_2 or nuclear norms for a sparse solution. ℓ_1 regularization is guaranteed to give a subspace-preserving affinity (i.e., there are no connections between points from different subspaces) under broad theoretical conditions, but the clusters may not be fully connected. ℓ_2 and nuclear norm regularization often improve connectivity, but give a subspace-preserving affinity only for independent subspaces. Mixed ℓ_1, ℓ_2 and nuclear norm regularization could offer a balance between the subspace-preserving and connectedness properties, but this comes at the cost of increased computational complexity. This paper focuses on using ℓ_1 norm and alleviating the corresponding connectivity problem by a simple yet efficient diffusion process on subspace affinity graphs. Without adding any tuning parameter , our method can achieve state-of-the-art clustering performance on Hopkins 155 and Extended Yale B data sets.

1. Introduction

Many computer vision problems, such as image compression [16], motion segmentation [26] and face clustering [15], deal with high-dimensional data. The high-dimensionality of the data not only increases the computational time and memory requirements of algorithms, but also decreases their performance due to the noise effect and insufficient number of samples with respect to the ambient space dimension, commonly referred to as the "curse of dimensionality" [3]. However, even though data are high-dimensional, their intrinsic dimension is often much smaller than the dimension of the ambient space, which has motivated the development of a number of techniques for finding a low-dimensional representation of a high-dimensional data set. Conventional techniques, such as Principal Component Analysis (PCA), assume that the data is drawn from a single low-dimensional subspace of the high-dimensional space. In practice, however, such high-dimensional data usually lie close to multiple low-dimensional subspaces corresponding to several classes or categories to which the data belong. In these scenarios, the task of clustering a high-dimensional data set into multiple classes becomes to the task of assigning each data point to its own subspace and recovering the underlying low-dimensional structure of the data, a problem known in the literature as *subspace clustering* [31].

In machine learning and computer vision communities, existing subspace clustering methods can be divided into four main categories, including algebraic methods [6, 32], iterative methods [30, 21], statistical methods [22, 27], and spectral clustering-based methods [39, 14, 4, 10, 19]. Among them, spectral-clustering based methods have become extremely popular due to their simplicity, theoretical correctness, and empirical success. These methods generally divide the problem into two steps: 1) Constructing an affinity matrix based on certain model and 2) applying spectral clustering to the affinity matrix. In this paper, we focus on the first step since the success of spectral clustering highly depends on having an appropriate affinity matrix.

State-of-the-art methods for constructing the affinity matrix in terms of subspace clustering are based on the *self-expressiveness property* of the data [9], i.e., each data point in a union of subspaces can be efficiently reconstructed by a linear combination of all other data points: $x_j = \sum_{i \neq j} c_{ij} x_i$, where the coefficient c_{ij} is used to define the affinity between points i and j as $w_{ij} = |c_{ij}| + |c_{ji}|$. However, this leads to an ill-posed problem with many possible solutions. To deal with this issue, the principle of *sparsity* is invoked. Specifically, every point is expressed as a sparse linear combination of all other data points by minimizing certain norm of coefficient matrix. This problem can then be written as:

$$\min_C ||C||, \quad s.t. \quad X = XC, \quad diag(C) = 0, \quad (1)$$

where $X = [x_1, ..., x_N]$ is the data matrix, $C =$

$[C_1, ..., C_N]$ is the coefficient matrix, $||\cdot||$ is a properly chosen regularizer.

The main difference among state-of-the-art methods lies in the choice of the regularizer. In Sparse Subspace Clustering (SSC) [9], the ℓ_1 norm is used for $||\cdot||$ as a convex relaxation of the ℓ_0 norm to promote the sparseness of C. While under broad theoretical conditions [10, 38] the sparse representation produced by SSC is guaranteed to be subspace-preserving (i.e.,$c_{ij} \neq 0$ only if x_i and x_j are in the same subspace), the affinity graph however may lack connectedness [] (i.e., data points from the same subspace may not form a connected component of the affinity graph due to the sparseness of the connections). In Low-Rank Representation (LRR) [19] and Low-Rank Subspace Clustering [11], the nuclear norm $||\cdot||_*$ is adopted as a convex relaxation of the rank function. One benefit of nuclear norm is that the coefficient matrix is generally dense, which alleviates the connectivity issue of sparse representation based methods. However, the representation matrix is known to be subspace preserving only when the subspaces are independent, which significantly limits its applicability.

In this paper, we propose to use ℓ_1 norm for sparsity constraint to best retain the subspace preserving property, and adopt a diffusion process on subspace of affinity graph to mitigate the connectedness problem. Specifically, the fist step is to learn a sparse affinity matrix by applying ℓ_1 norm on optimization problem of (1). Diffusion process is then adopted to spread the affinity values through the entire graph built upon the affinity matrix. Such a process is interpretable as a random walk on the graph, where a so-called transition (affinity) matrix defines probabilities for walking from one node to a neighboring node. One remarkable advantage here is, since the affinity matrix learned by ℓ_1 norm is subspace preserving, the random walk on this matrix is guaranteed to be subspace constrained. Therefore, the connectivity within subspace is significantly enhanced while the subspace preserving property remains unaltered. An illustrative example is given in Figure 1. It clearly demonstrates that based on the sparse affinity matrix obtained by ℓ_1 norm, the proposed method could evidently improve the affinity within subspaces while the sparsity between subspaces is retained, yielding affinity matrix with exactly block-diagonal structure.

2. Related Work

2.1. Subspace clustering

State-of-the-art subspace clustering methods are based on the self-expressiveness model. The main difference among those methods lies in the choice of the regularizer on the coefficient matrix. While different regularizers possess their own advantages and drawbacks, [34, 25, 37] propose to use mixed norms. For example, the low-rank sparse subspace clustering (LRSSC) method [34] combines ℓ_1 and nuclear norm regularizer:

$$||C||_* + \lambda||C||_1, \tag{2}$$

where λ controls the trade-off between the two regularizers. Likewise, [25, 37] propose to use a mixed ℓ_1 and ℓ_2 norm given by

$$\lambda||C||_1 + \frac{1-\lambda}{2}||C||_2, \tag{3}$$

where λ plays a trade-off role between the two norms. These methods basically attempt to bridge the gap between the subspace preserving and connectedness properties by the trade-off of different norms.

[12] proposes to explicitly impose a block-diagonal constraint by fixing the rank of Laplacian matrix. One benefit of this approach is that it can be applied to all the affinity construction methods straightforwardly. The Structured Sparse Subspace Clustering (SSSC) [18] integrates the two stages, affinity learning and spectral clustering, into one unified optimization framework. Their observation is that the clustering results can help the self-expressiveness model to yield a better affinity matrix.

2.2. Diffusion processes

Diffusion process is widely used in the field of retrieval [35, 1, 8, 33, 7], in which the task is retrieving the most similar instances to a provided query element from a potentially large database. Conventional approaches are usually based on analyzing pairwise affinity/distance values which are directly used to rank the most similar elements afterwards. Such approaches has the main limitation that the structure of the underlying data manifold is completely ignored. For this reason, instead of considering pairwise affinity individually, diffusion process is adopted to derive context sensitive measures, and it has shown to be an indispensable tool for improving retrieval performance [7].

Diffusion process is generally start with a affinity matrix $A_{N\times N}$. The first step is to interpret the matrix A as a undirected graph $G = (V, E)$, consisting of N nodes $v_i \in V$, and edges $e_{ij} \in E$ that link nodes to each other. The edge weights are fixed to the affinity values A_{ij}. Diffusion process then spreads the affinity values through the entire graph based on the defined edge weights.

Diffusion process can be interpreted as a Markov random walk on a graph $G = (V, E)$. To this end, we first define the transition matrix of random walk as

$$P = D^{-1}A, \tag{4}$$

where D is a diagonal matrix with $D_{ii} = \sum_{k=1}^{n} A(i,k)$. Obviously, P is a row-stochastic matrix, containing the transition probabilities for a random walk in the corresponding graph. With a simple undate rule $A_{t+1} = A_t P$, the

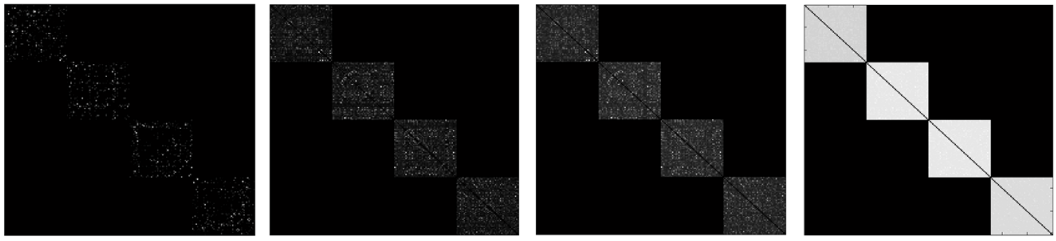

Figure 1. Visualization of affinity matrices obtained by different methods: (a) C_0 is the original coefficient matrix obtained by SSC(ℓ_1 norm), (b)$\backsim$(d) C_t are the matrices after t steps of diffusion process applied on C_0. Clearly, as the diffusion process goes on, the block-diagonal structure of the affinity matrix becomes evident.

affinity matrix after t steps of random walks can be obtained by

$$A_t = AP^t. \tag{5}$$

The random walk model was later extended to one of the most successful retrieval methods, the Google PageRank system [24]. The standard random walk is modified, and at each time step t a random walk step is done with probability α, whereas a random jump to an arbitrary node is made with probability $(1-\alpha)$. This leads to following update strategy

$$A_{t+1} = \alpha A_t P + (1-\alpha)Y, \tag{6}$$

where Y is probabilities of randomly jumping, which enables personalization for individual web user. Similar method was proposed in [40], namely Ranking on Manifolds. The different is the slightly adapted transition matrix $P = D^{-1/2}AD^{-1/2}$.

The drawback of the diffusion process mentioned above is that the process is applied to the entire graph, which can be heavily influenced by noisy edges. Current state-of-the-art methods [36] restricted the diffusion process to the K nearest neighbor (KNN) graph. Given affinity matrix A, this method sets $A_{ij} = A_{ij}$ only if $x_j \in KNN(x_i)$, otherwise $A_{ij} = 0$, and comes up with a new update strategy as

$$A_{t+1} = PA_tP^T. \tag{7}$$

While this approach consistently yields state-of-the-art performance in the field of retrieval, it is observed that in practice one needs to set neighborhood K manually and the performance is sensitive to the choice of K [7].

Paper Contributions. In this paper, we exploit diffusion process to mitigate connectedness problem of ℓ_1 norm in terms of subspace clustering. To the best of our knowledge, this is the first attempt to adopt the idea of diffusion to this field. Since the idea of using ℓ_1 norm for subspace clustering is originally from Sparse Subspace Clustering (SSC), we refer our method as Diffusion-based Sparse Subspace Clustering (DSSC). Our main contributions can be summarized as:

1. For subspace clustering, instead of adding different norms to balance the subspace preserving and connectivity properties, we come up with using ℓ_1 norm and alleviating the corresponding connectivity problem by a simple yet efficient diffusion process. Without adding any tuning parameter, the widely existing gap between the two properties is well bridged.

2. From the diffusion point of view, we show that instead of choosing KNN neighbors based on Euclidean distance, the sparse property of ℓ_1 norm provides manifold-aware neighborhood construction, which are locally constrained in corresponding subspaces. Moreover, the tuning parameter k of original diffusion is eliminated.

3. We present experiments on both synthetic data sets and real computer vision data sets that demonstrate the superiority of the proposed method compared to other state-of-the-art methods.

3. Diffusion based Sparse Subspace Clustering

Before introducing the proposed approach, we first formulate the addressed problem in this paper.

Problem 1 *Given a collection of data points* $\{x_j \in \mathbb{R}^D\}_{j=1}^N$ *drawn from an unknown union of k subspaces* $\{S_i\}_{i=1}^k$ *of unknown dimensions* $d_i = dim(S_i), 0 < d_i < D, i = 1, ..., k$. *The goal is to segment these points into their corresponding subspaces.*

Sparse Subspace Clustering (SSC) attempts to solve the problem based on the so-called self-expressiveness model, which states that each data point can be expressed as a linear combination of all other data points, i.e., $X = XC + E$, where C is the coefficient matrix and E is the matrix of error (noises or outliers). In principle, this leads to an ill-posed problem with many possible solutions. Thus, the

sparsity constraint is invoked by ℓ_1 minimization, leading to the following optimization problem

$$\min_C ||C||_1 + ||E||_\ell \quad s.t. \quad X = XC + E, \quad diag(C) = 0, \tag{8}$$

where the Frobenius norm or ℓ_1 norm is used for $||\cdot||_\ell$ to handle noise or outliers. The SSC algorithm proceeds by solving the optimization problem in (8) using the ADMM method. The optimal coefficient C is then used to define an affinity matrix $|C| + |C^T|$. The segmentation of the data is finally obtained by applying spectral clustering to the normalized Laplacian.

While SSC works well in practice, one possible drawback is ℓ_1 finds the sparse representation of each data point individually. In the case of clean data, SSC is guaranteed to be subspace-preserving, i.e., there are no connections between points from different subspaces. However, the within-subspace connections are usually sparse, i.e., C_{ij} could be zero even x_i and x_j are in the same subspace. It is not a problem as long as the connections are still subspace-preserving. But in the case of noisy data, there is no theoretical guarantee that the nonzero coefficients correspond to points in the same subspace. Imaging there are connections between points from different subspaces, spectral clustering cannot be able to appropriately cut the graph, as shown in Figure 3.

Diffusion process is capable to deal with this issue by exploiting the contextual affinities. Given the affinity matrix $W = |C| + |C^T|$, where C is obtained by solving (8), diffusion process is encoded into computing the power of the affinity matrix, which is

$$W_t = W^t, \tag{9}$$

where t corresponds t steps of diffusion process. Obviously, such a process is sensitive to the step t. In order to make the diffusion process independent from t, we consider the accumulation of all t. Thus, the diffusion process is

$$W_t = \sum_{i=0}^{t} W^i. \tag{10}$$

We assume that W is nonnegative and the sum of each row is smaller than one. A matrix W that satisfies these requirements can be easily constructed from a stochastic matrix. Note that the absolute values of the eigenvalues is bounded by the maximum of the rowwise sums. Therefore, the maximum of the absolute values of the eigenvalues of W is smaller than one. Consequently, (10) converges to a fixed and nontrivial solution given by $\lim_{t\to\infty} W_t = (I - W)^{-1}$, where I is the identity matrix.

To further incorporate the contextual affinity, it is shown in [36] that the diffusion process on higher order tensor product graph is promising for revealing the intrinsic relation between data points. Given graph $G = (V, E)$ constructed from affinity matrix W, the tensor product graph is defined as the Kronecker product of original graph, $\mathbb{G} = G \bigotimes G$. The corresponding affinity matrix is $\mathbb{W} = W \bigotimes W$. In particular, we have

$$\mathbb{W}(\alpha, \beta, i, j) = W(\alpha, \beta) \cdot W(i, j) = w_{\alpha,\beta} \cdot w_{i.j}. \tag{11}$$

Thus, if $W \in \mathcal{R}^{n\times n}$, then $\mathbb{W} = W \bigotimes W \in \mathcal{R}^{nn\times nn}$.

The diffusion process is then defined on the higher order tensor as

$$\mathbb{W}_t = \sum_{i=0}^{t} \mathbb{W}^i. \tag{12}$$

Since the sum of each row of W is smaller than 1, we have

$$\sum_{\beta j} \mathbb{W}(\alpha, \beta, i, j) = \sum_{\beta j} w_{\alpha\beta} w_{ij} = \sum_{\beta} w_{\alpha\beta} \sum_{j} w_{ij} < 1. \tag{13}$$

As is the case for (10), the process (12) also converges to a fixed and nontrivial solution,

$$\lim_{t\to\infty} \mathbb{W}_t = \lim_{t\to\infty} \sum_{i=0}^{t} \mathbb{W}^i = (I - \mathbb{W})^{-1} \tag{14}$$

Since our goal is to learn a new affinity matrix W^* of size $n \times n$, it is defined as

$$W^* = vec^{-1}((I - \mathbb{W})^{-1} vec(I)), \tag{15}$$

where I is the identity matrix and vec is an operator that stacks the columns of a matrix one after the next into a column vector. The inverse of vec is denoted as vec^{-1}.

While the tensor graph provides adequate underlying structure of the data, it is impractical for large scale problems due to the demand of high storage and computing cost. Therefore, we use an iterative algorithm for the diffusion process on tensor graph. First, we define $A_1 = W$ and the update strategy as

$$A_{t+1} = W A_t W^T + I, \tag{16}$$

where I is the identity matrix. The diffusion process becomes the iteration of (16) until convergence. To prove the convergence of (16), we first transform (16) to

$$
\begin{aligned}
A_{t+1} &= WA_tW^T + I = W(WA_{t-1}W^T + I)W^T + I \\
&= W^2A_{t-1}(W^T)^2 + WIW^T + I = ... \\
&= W^tW(W^T)^t + W^{t-1}I(W^T)^{t-1} + ... + I \\
&= W^tW(W^T)^t + \sum_{i=0}^{t-1} W^iI(W^T)^i. \qquad (17)
\end{aligned}
$$

Since we assume the sum of each row of $W < 1$, we have $\lim_{t\to\infty} W^tW(W^T)t = 0$, consequently,

$$
\lim_{t\to\infty} A_{t+1} = \lim_{t\to\infty} \sum_{i=0}^{t-1} W^iI(W^T)^i. \qquad (18)
$$

As shown in [36], it can be proven by induction that

$$
\lim_{t\to\infty} \sum_{i=0}^{t-1} W^iI(W^T)^i = vec^{-1}((I - \mathbb{W})^{-1}vec(I)). \quad (19)
$$

So we have

$$
\lim_{t\to\infty} A_{t+1} = vec^{-1}((I - \mathbb{W})^{-1}vec(I)). \qquad (20)
$$

Hence, the iterative algorithm (16) yields the same affinities as the diffusion process on tensor graph.

3.1. A graph view of the diffusion process

With ℓ_1 minimization, SSC seeks sparse representation for each data point individually. The corresponding affinity graph reveals the pairwise affinity as the "shortest path" between them, which is susceptible to noise. The proposed DSSC derives the affinity by considering the "volume of paths" through a diffusion process. One remarkable invention is, with ℓ_1 norm, the "volume of paths" are restricted in subspaces, resulting in the enhancement of within-subspace affinity and constant of between-subspace affinity.

Figure 2 shows that in the case of clean data, SSC is guaranteed to be subspace-preserving, i.e., there are no connections between subspaces. Thus, spectral clustering can properly cut the graph. However, due to the sparse property of ℓ_1 norm, it is very likely that two points in the same subspace are not connected. While in such a case, the diffusion process can accurately complete the graph for each subspace, leading to more robust subspace graphs.

In the case of noisy data, the advantage of diffusion becomes more significant. As shown in Figure 3, when there are noisy edges between subspaces, spectral clustering may not be able to find the ideal cut of the graph. As the pairwise affinity is individually computed in SSC, it may be difficult to distinguish "noisy edge" with "real edge". While in DSSC, due to the ability of combining contextual information, the within-subspace affinities are evidently enhanced

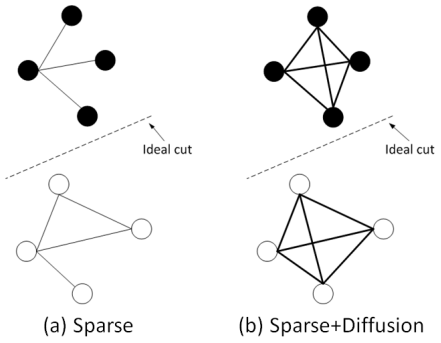

Figure 2. An example of affinity graph for clean data, obtained by different methods: (a) sparse, (b) sparse with diffusion.

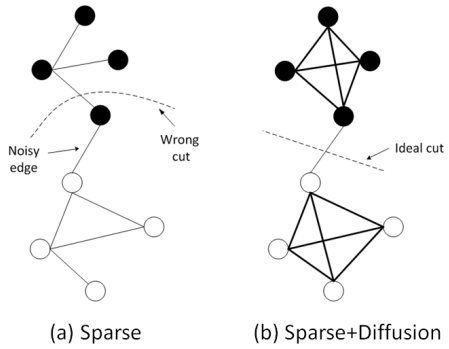

Figure 3. An example of affinity graph for noisy data, obtained by different methods: (a) sparse, (b) sparse with diffusion.

so that the noisy edges can be properly cut by spectral clustering. It should be noticed that real world data are usually noisy, which explains the significant improvements on real world data sets, shown in Section 4.2 and Section 4.3.

3.2. A random walk view of the diffusion process

As we mentioned, the proposed diffusion process can be interpreted as Markov random walks [17], where the transition (affinity) matrix defines probabilities for walking from one node to a neighboring node. The crucial observation is that considering a affinity graph, there will be many connections within a subspace, and fewer connections between subspaces. Therefore, if we start a walker at one node and then randomly travel to a connected node, the random walker is more likely to stay within the subspace than travel between different subspaces. Intuitively, diffusion process can be imaged as random walks starting at every node on the graph. As the random walks go on, the probabilities of traveling between nodes within subspace will increase, while the probabilities between subspaces decrease.

Given affinity matrix A, one can easily build a stochastic transition matrix P by (4). The probability distribution of the positions of a random walk, starting at node j, is given

by the j-th row of the transition matrix. After t time steps, the probability of a random walk starting at node j at time 0, to be at node i at time t is thus,

$$p_t(i|j) = P^t_{i,j}. \tag{21}$$

Under mild conditions [5], we obtain in this way an ergodic Markov process with a single stationary distribution $\pi = (\pi_1, ..., \pi_n)$, where $\pi_i = d_i/vol(V)$, d_i is the degree of node i, V is the set of all nodes. It is easy to verify that this distribution is a right-eigenvector of the t-step transition matrix (for every t), e.g. $P^T\pi = \pi$, since $\pi_i = \sum P_{i,j}\pi_j$.

As we know, there is a tight relation between random walks and spectral segmentation [23]. We take the Normalized Cut [28] (NCut) as an example of spectral segmentation. The goal of NCut algorithm is to segment an image into two disjoint parts by minimizing

$$NCut(A, \overline{A}) = \left(\frac{1}{vol(A)} + \frac{1}{vol(\overline{A})}\right) \sum_{i\in A, j\in\overline{A}} A_{ij}. \tag{22}$$

As derived by [23], the objective function of NCut (27) can be expressed in the framework of random walk as follows. For two disjoint subsets A, $B \subset V$, assume we run a random walk starting with X_0 in the stationary distribution π. We define

$$P(B|A) = P(X_1 \in B|X_0 \in A) \tag{23}$$

as the probability of the random walk transition from set A to set B. First of all it can be observed that

$$\begin{aligned} P(X_0 \in A, X_1 \in B) &= \sum_{i\in A, j\in B} P(X_i, X_j) \\ &= \sum_{i\in A, j\in B} \pi_i p_{ij} \\ &= \sum_{i\in A, j\in B} \frac{d_i}{vol(V)} \frac{a_{ij}}{d_i} \\ &= \frac{1}{vol(V)} \sum_{i\in A, j\in B} a_{ij}. \end{aligned} \tag{24}$$

With the Bayes' Rules, we have the posterior probability

$$\begin{aligned} P(B|A) &= \frac{P(X_0 \in A, X_1 \in B)}{P(X_0 \in A} \\ &= \left(\frac{1}{vol(V)} \sum_{i\in A, j\in B} a_{ij}\right) \left(\frac{vol(A)}{vol(V)}\right)^{-1} \\ &= \frac{1}{vol(A)} \sum_{i\in A, j\in B} a_{ij}. \end{aligned} \tag{25}$$

Follow this, the definition of NCut can be written as

$$NCut(A, \overline{A}) = P(\overline{A}|A) + P(A|\overline{A}). \tag{26}$$

While originally the objective of NCut algorithm is to find those regions that the between connections are minimized, it can be interpreted as to find regions in which the probabilities of random walkers escape from these regions are low, with the theory of random walk.

Denoting the transition matrix after t time steps random walk as P_t, we know that $P_t(\overline{A}|A) < P(\overline{A}|A)$ while $P_t(A|A) > P(A|A)$. Let $NCut^*$ be the NCut criterion on the transition matrix after random walk, the following holds,

$$NCut^*(A, \overline{A}) = P_t(\overline{A}|A) + P_t(A|\overline{A}) < NCut(A, \overline{A}), \tag{27}$$

which explains why random walk (diffusion process) can improve the performance of spectral segmentation.

4. Experiments

Experiments are demonstrated in this section. We evaluate the proposed DSSC approach on a synthetic data set, a motion segmentation data set, and a face clustering data set to validate its effectiveness.

Experimental Setup. Since the proposed DSSC is built upon the standard SSC [10], we keep all settings in DSSC the same as in SSC. As for diffusion process, since it is guaranteed to converge, we set the iteration to 200 for all experiments. The performance is validated by clustering error, which is measured by

$$clustering\ \ error = \frac{\#\ of\ missclassified\ \ points}{total\ \#\ of\ points}. \tag{28}$$

4.1. Experiments on synthetic data

Data Generation. We construct 5 subspaces $\{S\}_{i=1}^5 \subset \mathbb{R}^{100}$ whose bases $\{U_i\}_{i=1}^5$ are obtained by $U_i = T_i U, 1 \leq i \leq 5$, where U is a random matrix of dimension 100×5 and $T_i^{100\times 100}$ is a random rotation. We sample 50 data points from each subspace by $X_i = U_i Q_i$, where the entries of $Q_i \in \mathbb{R}^{5,50}$ are i.i.d. samples from a standard Gaussian. Some data vectors x are then randomly chosen to corrupt by Gaussian noise with zero mean and variance $0.3||x||$.

Experimental results are presented in Table 1. It can be observed that the proposed DSSC consistently outperforms SSC. In the case of clean data (no corruptions), both SSC and DSSC achieve perfect clustering with error rates equal to 0. As the corruptions increase, so do the chance of adding noise edges into the corresponding affinity graph. In this case, spectral clustering cannot find the appropriate cut of

Table 1. Clustering Errors on Synthetic Dataset.

Corruptions	0	10%	20%	30%	40%	50%	60%	70%	80%	90%	100%
SSC	0	2.88	6.90	13.26	18.66	24.62	30.16	34.42	38.80	40.20	43.00
DSSC	**0**	**1.60**	**4.80**	**9.74**	**14.40**	**17.64**	**22.80**	**28.98**	**32.80**	**36.00**	**39.74**

Figure 4. Example frames from videos in the Hopkins 155 [29].

the graph in SSC. While in DSSC, those noisy edges are well detected due to the within subspaces edges are significantly enhanced by the diffusion process, as demonstrated in Figure 2 and Figure 3.

4.2. Experiments on Motion Segmentation

Motion segmentation refers to the problem of segmenting a video sequence of multiple rigidly moving objects into multiple spatiotemporal regions that correspond to different motions in the scene (see Figure 4). This problem is often solved by extracting and tracking a set of feature points through all frames of the video. Each data point, which is also called a feature trajectory, corresponds to a vector obtained by stacking all feature points. Under the affine projection model, all feature trajectories associated with a single rigid motion lie in an affine subspace of dimension at most 3 [10]. Therefore, motion segmentation reduces to clustering of these trajectories in a union of subspaces.

We evaluate the proposed DSSC algorithm with other state-of-the-art subspace clustering methods, i.e., SSC [10], SSSC [18], LRR [19], LSR [20], BDSSC [12], LRSC [11], on the Hopkins 155 motion segmentation data set [29] for the multi-view affine motion segmentation. It consists of 155 video sequences, where 120 of the videos have two motions and 35 of the videos have three motions. We evaluate average performance on three cases: 2 motions, 3 motions, and all. Experimental results are presented in Table 2. The result for LRSC is cited from [11], while the others are cited from [18]. Note that the proposed DSSC achieves the best performances on all cases.

4.3. Experiments on Face Clustering

Given face images of multiple subjects acquired with a fixed pose and varying illumination, we consider the problem of clustering images according to their subjects. It has been shown that, under the Lambertian assumption, images of a subject with a fixed pose and varying illumination lie close to a linear subspace of dimension 9 [2]. Thus, face clustering can be also considered as a subspace clustering problem, where each subject lies in a 9D subspace.

We evaluate the clustering performance of the proposed DSSC as well as other state-of-the-art methods on the Extended Yale B data set [13]. It contains 2,414 frontal face images of 38 subjects, with 64 images per subject acquired under different illumination conditions. In our experiments, we follow the same settings introduced in [10]. It should be noticed that the Extended Yale B data set is more challenging for subspace segmentation than the Hopkins 155 data set due to the heave noise, high-dimensional space, and large number of subspace in the data.

Experimental results are presented in Table 3. The results are directly cited from [18], which is a fair comparison as we use exactly same experimental settings. It can be observed that the proposed DSSC performs the best results on all cases. Note that on this more challenging data set, DSSC achieves significant improvements compared to the state-of-the-art methods. It is worth mentioning that these significant improvements are achieved by a parameter-free diffusion process, while other methods usually add in tuning parameters for flexibility.

5. Conclusion

In this work, we investigated diffusion process for sparse subspace clustering and proposed a new subspace clustering method, namely DSSC. Specifically, after we obtained affinity matrix by ℓ_1 minimization, we adopted a diffusion process to spread the affinity value. With the subspace-preserving property of ℓ_1 norm, such a diffusion process is remarkably constrained within each subspaces, yielding enhanced within-subspace affinity and unaltered between-subspace affinity. Moreover, we explained the diffusion process in the views of graph and random walk, and gave theoretical justifications on how does the diffusion process improve the performance of spectral clustering. Extensive experiments verified that our proposed diffusion based sparse subspace clustering method, without adding in tuning parameter, could significantly improve the state-of-the-art performance.

References

[1] X. Bai, X. Yang, L. J. Latecki, W. Liu, and Z. Tu. Learning context-sensitive shape similarity by graph transduction. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(5):861–874, 2010.

[2] R. Basri and D. W. Jacobs. Lambertian reflectance and linear subspaces. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 25(2):218–233, 2003.

Table 2. Motion Segmentation Errors(%) on Hopkins 155 Dataset.

Methods	LRR [19]	LRSC [11]	LSR [20]	BDSSC [12]	SSC [10]	SSSC[18]	DSSC
2 motions	3.76	3.69	2.20	2.29	1.95	1.94	**1.68**
3 motions	9.92	7.69	7.13	4.95	4.94	4.92	**4.64**
All	5.15	4.59	3.31	2.89	2.63	2.61	**2.35**

Table 3. Clustering Errors (%) on the Extended Yale B Dataset.

Methods	LRR [19]	LRSC [11]	LSR [20]	BDSSC [12]	SSC [10]	SSSC[18]	DSSC
2 subjects	6.74	3.15	6.72	3.90	1.87	1.27	**0.61**
3 subjects	9.30	4.71	9.25	17.70	3.35	2.71	**1.25**
5 subjects	13.94	13.06	13.87	27.50	4.32	3.41	**2.80**
8 subjects	25.61	21.25	25.98	33.20	5.99	4.15	**4.04**
10 subjects	29.53	29.58	28.33	39.53	7.29	5.16	**4.84**

[3] R. E. Bellman. *Dynamic Programming*. Princeton Univ. Press, 1957.

[4] G. Chen and G. Lerman. Spectral curvature clustering (SCC). *International Journal of Computer Vision*, 81(3):317–330, 2009.

[5] F. R. Chung. *Spectral graph theory*, volume 92. American Mathematical Soc., 1997.

[6] J. P. Costeira and T. Kanade. A multibody factorization method for independently moving objects. *International Journal of Computer Vision*, 29(3):159–179, 1998.

[7] M. Donoser and H. Bischof. Diffusion processes for retrieval revisited. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1320–1327, 2013.

[8] A. Egozi, Y. Keller, and H. Guterman. Improving shape retrieval by spectral matching and meta similarity. *Image Processing, IEEE Transactions on*, 19(5):1319–1327, 2010.

[9] E. Elhamifar and R. Vidal. Sparse subspace clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2790–2797, 2009.

[10] E. Elhamifar and R. Vidal. Sparse subspace clustering: Algorithm, theory, and applications. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 35(11):2765–2781, 2013.

[11] P. Favaro, R. Vidal, and A. Ravichandran. A closed form solution to robust subspace estimation and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1801–1807, 2011.

[12] J. Feng, Z. Lin, H. Xu, and S. Yan. Robust subspace segmentation with block-diagonal prior. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 3818–3825, 2014.

[13] A. S. Georghiades, P. N. Belhumeur, and D. J. Kriegman. From few to many: Illumination cone models for face recognition under variable lighting and pose. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 23(6):643–660, 2001.

[14] A. Goh and R. Vidal. Segmenting motions of different types by unsupervised manifold clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–6, 2007.

[15] J. Ho, M.-H. Yang, J. Lim, K.-C. Lee, and D. Kriegman. Clustering appearances of objects under varying illumination conditions. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–11, 2003.

[16] W. Hong, J. Wright, K. Huang, and Y. Ma. Multiscale hybrid linear models for lossy image representation. *Image Processing, IEEE Transactions on*, 15(12):3655–3671, 2006.

[17] M. S. T. Jaakkola and M. Szummer. Partially labeled classification with markov random walks. *Advances in neural information processing systems*, 14:945–952, 2002.

[18] C.-G. Li and R. Vidal. Structured sparse subspace clustering: A unified optimization framework. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 277–286, 2015.

[19] G. Liu, Z. Lin, and Y. Yu. Robust subspace segmentation by low-rank representation. In *International Conference on Machine Learning*, pages 663–670, 2010.

[20] C.-Y. Lu, H. Min, Z.-Q. Zhao, L. Zhu, D.-S. Huang, and S. Yan. Robust and efficient subspace segmentation via least squares regression. In *European Conference on Computer Vision*, pages 347–360. 2012.

[21] L. Lu and R. Vidal. Combined central and subspace clustering for computer vision applications. In *International Conference on Machine Learning*, pages 593–600, 2006.

[22] Y. Ma, H. Derksen, W. Hong, and J. Wright. Segmentation of multivariate mixed data via lossy data coding and compression. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, (9):1546–1562, 2007.

[23] M. Meila and J. Shi. A random walks view of spectral segmentation. 2001.

[24] L. Page, S. Brin, R. Motwani, and T. Winograd. The PageRank citation ranking: bringing order to the web. 1999.

[25] Y. Panagakis and C. Kotropoulos. Elastic net subspace clustering applied to pop/rock music structure analysis. *Pattern Recognition Letters*, 38:46–53, 2014.

[26] S. Rao, R. Tron, R. Vidal, and Y. Ma. Motion segmentation in the presence of outlying, incomplete, or corrupted trajectories. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 32(10):1832–1845, 2010.

[27] S. R. Rao, R. Tron, R. Vidal, and Y. Ma. Motion segmentation via robust subspace separation in the presence of outlying, incomplete, or corrupted trajectories. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2008.

[28] J. Shi and J. Malik. Normalized cuts and image segmentation. *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 22(8):888–905, 2000.

[29] R. Tron and R. Vidal. A benchmark for the comparison of 3-d motion segmentation algorithms. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 1–8, 2007.

[30] P. Tseng. Nearest q-flat to m points. *Journal of Optimization Theory and Applications*, 105(1):249–252, 2000.

[31] R. Vidal. A tutorial on subspace clustering. *IEEE Signal Processing Magazine*, 28(2):52–68, 2010.

[32] R. Vidal, Y. Ma, and S. Sastry. Generalized principal component analysis (GPCA). *Pattern Analysis and Machine Intelligence, IEEE Transactions on*, 27(12):1945–1959, 2005.

[33] B. Wang and Z. Tu. Affinity learning via self-diffusion for image segmentation and clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2312–2319, 2012.

[34] Y.-X. Wang, H. Xu, and C. Leng. Provable subspace clustering: When lrr meets ssc. In *Advances in Neural Information Processing Systems*, pages 64–72, 2013.

[35] X. Yang, S. Koknar-Tezel, and L. J. Latecki. Locally constrained diffusion process on locally densified distance spaces with applications to shape retrieval. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 357–364, 2009.

[36] X. Yang and L. J. Latecki. Affinity learning on a tensor product graph with applications to shape and image retrieval. In *IEEE Conference on Computer Vision and Pattern Recognition*, pages 2369–2376, 2011.

[37] C. You, C.-G. Li, D. Robinson, and R. Vidal. Oracle based active set algorithm for scalable elastic net subspace clustering. In *IEEE Conference on Computer Vision and Pattern Recognition*, 2016.

[38] C. You and R. Vidal. Geometric conditions for subspace-sparse recovery. In *International Conference on Machine Learning*, pages 1585–1593, 2015.

[39] T. Zhang, A. Szlam, Y. Wang, and G. Lerman. Hybrid linear modeling via local best-fit flats. *International Journal of Computer Vision*, 100(3):217–240, 2012.

[40] D. Zhou, J. Weston, A. Gretton, O. Bousquet, and B. Schölkopf. Ranking on data manifolds. *Advances in neural information processing systems*, 16:169–176, 2004.